IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF APPEALS

Appeal No. _____

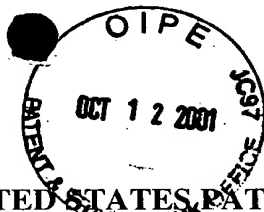In re Application of: JOHN C. THACKER ET AL.

Serial No.: 09/356,997

Filed: July 20, 1999

For: METHOD AND APPARATUS FOR INTERNET CACHE CONTENT
DELIVERY VIA A DATA DISTRIBUTION SYSTEM

APPELLANTS' BRIEF ON APPEAL

Kenneth W. Float
The Law Offices of Kenneth W. Float
Office address: 2 Shire, Coto de Caza, CA 92679
Mailing address: P.O. Box 80790, Rancho Santa Margarita, CA 92688

APPELLANT'S BRIEF ON APPEAL

RECEIVED

IOCT 1 6 2001

Group 2100

Commissioner of Patents and Trademarks
Washington, D. C. 20231

Sir:

This is Appellants' brief on appeal from the decision of the Examiner in the Office Action dated April 23, 2001 finally rejecting Claims 1-19 pending in the above-identified patent application. This brief is submitted in accordance with the provisions of 37 C.F.R. §1.192.

REAL PARTY IN INTEREST

The real party in interest is CyberStar, L.P. which acquired rights to the present application by way of an assignment from the inventors.

RELATED APPEALS AND INTERFERENCES

There are no other appeals or interferences known to appellants, appellant's legal representative, or the assignee which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

STATUS OF CLAIMS

Claims 1-19 are currently pending in this application and were finally rejected in the Office Action dated April 23, 2001. Appellants 'appeals from this final rejection.

10/16/2001 GTEFFERA 00000048 09356997

01 FC:120   320.00 OP   **STATUS OF AMENDMENTS**

With regard to the status of amendments, three Office Actions were issued during prosecution of this application. Amendments were made to the Claims in response to the first and second Office Actions. No amendments were made in response to the final Office Action dated April 23, 2001. The Claims as they currently stand are presented in the Appendix.

SUMMARY OF INVENTION

The present invention provides for a cache system (10) and caching method (50) that is used to improve data delivery service provided by a data distribution system (30), such as a satellite-based data distribution system (30). The present invention uses a master cache (18) in

which predictive harvesting of content is performed based upon probability distributions of individual caches satisfying requests from their users (clients). Content in the master cache (18) is distributed to user sites by way of the satellite-based data distribution system (30), and local caches (20) at user sites are automatically loaded with incoming content.

An exemplary system (10) comprises a master cache (18) that receives content for distribution by the data distribution system (30) to one or more users. A gateway (15) receives content distributed by the data distribution system (30) from the master cache (18). One or more local caches (20) store the content received by the gateway (15) destined for the one or more users. Harvesting software (17) processes information from the master cache (18) and the gateway (15) that corresponds to probability distributions that the local caches (20) satisfy requests from their respective users to predictively distribute the desired content to the respective users.

Content distributed by the data distribution system (30) to the gateway (15) is transferred into the local cache (20) by creating a pseudo client (41) on the gateway (15), receiving an interrupt at the pseudo client (41) indicating that that content has arrived at the gateway (15), enabling the gateway (15) as a sibling cache (42) for the local cache (20), requesting content to be transferred from the sibling cache (42) to the local cache (20), verifying that content has been transferred to the local cache (20) during the transfer process, and disabling the gateway (15) as a sibling cache (42) of the local cache (20) at the completion of the process. The local cache (20) retrieved the content from the sibling cache (42) until all content has been transferred.

The master cache (18) is built by processing statistics derived from the master cache (18) and the local cache (20) to produces a list (62) of content to add to the master cache (15) and a list (64) of content to delete from the master cache (18). A pseudo client (20) is formed that retrieves and verifies the content to be added to the master cache (18). The verified content is then transmitted from the master cache (18) to the local cache (20).

ISSUES

The sole issue in this appeal is whether Claims 1-16 are unpatentable under 35 U.S.C. § 103(a) over U.S. Patent No. 6,085,193 issued to Malkin et al.

GROUPING OF CLAIMS

With regard to the specific grounds of rejection that are in issue, it is respectfully submitted that Claims 1-8, Claim 9, and Claims 10-12 should be considered separately.

ARGUMENT

The issue in this appeal is whether Claims 1-16 are unpatentable under 35 U.S.C. § 103(a) over U.S. Patent No. 6,085,193 issued to Malkin et al. It is respectfully submitted that the Examiner's rejection of Claims 11-16 is in error. As will be argued below, it is also respect-

fully submitted that the Examiner has extended the teachings of the Malkin et al. patent and has used hindsight reconstruction and "Official Notice" to reject the present invention.

The Examiner's position is that the Malkin et al. patent discloses "harvesting software coupled to the master cache and the uplink for processing information corresponding to probability distributions that the local caches satisfy requests from their respective users to predictively distribute the desired content to the respective users (see fig. 2; col. 6, Malkin discloses proxy server logic 295 which predictively retrieves content for the client from the content server)." It is respectfully submitted that the Malkin et al. patent does not disclose or suggest this aspect of the present invention, and particularly not without the use of hindsight reconstruction.

The Examiner admitted that "Malkin does not explicitly disclose a gateway connected to the master cache. Malkin discloses that proxy servers are used to retrieve content from the content server (master cache) (see col. 4-6)." However, in order to get around this deficiency, the Examiner has taken "Official Notice that the concept and advantages of using a Gateway to connect a resource distribution network to a subscriber or client network is old and well known in the network communication art."

Claim 1 calls for a caching system for use with a data distribution system, comprising:

a master cache for receiving content for distribution by the data distribution system to one or more users;
a gateway for receiving content that is distributed by the data distribution system from the master cache;
one or more local caches for storing the content received by the gateway destined for the one or more users; and
harvesting software coupled to the master cache and the gateway for processing information corresponding to probability distributions that the local caches satisfy requests from their respective users to predictively distribute the desired content to the respective users.

It is respectfully submitted that Malkin et al. patent does not require, nor would one skilled in the art be inclined to use, a gateway for receiving content that is distributed by the data distribution system from the master cache as is presently claimed. The Malkin et al. system does not use a gateway, and this is extraneous and not required by the Malkin et al. system. Therefore, it is respectfully submitted that one skilled in the art would not employ a gateway in the Malkin et al. system. This usage has been conjured up by the Examiner and is clearly based upon hindsight reconstruction using the teachings of the Malkin et al. patent in light of Applicants' own teachings. The only disclosure or suggestion regarding use of a gateway is found in the present specification. The gateway is used in the preferred embodiment of the present invention to receive content that is distributed by the data distribution system, and in particular a satellite-based data distribution system. A satellite-based data distribution system is not disclosed or suggested in the Malkin et al. patent.

While Examiner has taken Official Notice that the concept and advantages of using a gateway to connect a resource distribution network to a subscriber or client network is old and

well known in the network communication art, it is respectfully submitted that there is no disclosure or suggestion in the Malkin et al. patent as to the need for a gateway or where such a gateway would be used. It is respectfully submitted that a gateway is not disclosed or suggested in the Malkin et al. patent because it is not necessary since the Internet links the clients to the content server. Furthermore, if, as the Examiner states, gateways are old and well known in the network communication art, then if one were necessary in the Malkin et al. system, then it would have been disclosed or suggested by them. This clearly provides an indication that a gateway is not desired or required in the Malkin et al. system, and that inclusion of one using the assertions of the Examiner amount to hindsight reconstruction.

Furthermore, as has been argued in prior responses, it is respectfully submitted that Malkin et al. patent does not disclose or suggest the presently claimed "harvesting software ... for processing information corresponding to probability distributions that the local caches satisfy requests from their respective users to predictively distribute the desired content to the respective users." There is no disclosure or suggestion regarding "probability distributions" or " processing information corresponding to probability distributions that the local caches satisfy requests from their respective users" or software that "predictively distribute the desired content to the respective users" as is presently claimed.

The Examiner asserted that Fig. 2 and column 6 of the Malkin et al. patent teaches the presently claimed harvesting software. Clearly, this conclusion is in error. The Examiner admitted that the Malkin et al. patent fails to disclose the use of a gateway, which the Examiner has added to their system by way of "Official Notice". Notwithstanding this, the Examiner asserts that the Malkin et al. patent discloses "harvesting software coupled to the master cache and the gateway". This conclusion can only be reached based upon the use of hindsight reconstruction. The Malkin et al. patent does not disclose or suggest that harvesting software is disposed on both a master cache and a gateway. This is only taught in the present specification.

Furthermore, the Examiner stated that the Malkin et al. patent teaches "one or more local caches for storing the content received by the gateway destined for the one or more users (see fig. 1-2; col. 7, Malkin discloses second prefetch cache 249 at the client computers)". It is respectfully submitted that the present invention has "one or more local caches for storing the content received by the gateway". The Malkin et al. patent does not disclose or suggest the use of a gateway, and this is not disclosed or suggested without the use of hindsight reconstruction.

While the Malkin et al. patent discloses the use of statistics tables, it is respectfully submitted that there is no disclosure or suggestion "harvesting software coupled to the master cache and the gateway for processing information corresponding to probability distributions that the local caches satisfy requests from their respective users to predictively distribute the desired content to the respective users." The Malkin et al. patent states that it discloses a "method for prefetching data identifies data access patterns and prefetches select information based on a dynamic interpretation of the data access patterns". It is respectfully submitted that this is not a disclosure or suggestion regarding processing information corresponding to

probability distributions to predictively distribute the desired content to the respective users. There is no predictive distribution disclosed or suggested in the Malkin et al. patent. The term "predict" is used in the Malkin et al. patent one time to discuss dynamically tracking of viewer reference behavior, static analysis or mining on Web logs. The mere use of statistics tables is not a disclosure or suggestion regarding "processing information corresponding to probability distributions that the local caches satisfy requests from their respective users to predictively distribute the desired content to the respective users", as is presently claimed.

Thusly, there is no disclosure or suggestion in the Malkin et al. patent regarding predictive distribution of desired content to the respective users. The Malkin et al. patent contains no discussion regarding predictive distribution of information to users. The Malkin et al. patent does not disclose or suggest anything regarding harvesting software that processes information contained in transmit hit/miss data and probability tables. Probability tables are not disclosed or suggested in the Malkin et al. patent.

Therefore, with specific regard to Claim 1, it is therefore respectfully submitted that the Malkin et al. patent does not disclose or suggest "a gateway for receiving content that is distributed by the data distribution system from the master cache" or "harvesting software coupled to the master cache and the gateway for processing information corresponding to probability distributions that the local caches satisfy requests from their respective users to predictively distribute the desired content to the respective users", as is recited therein.

Accordingly, it is respectfully submitted that Claim 1 is not obvious in view of the Malkin et al. patent and is patentable thereover. Therefore, reversal of the Examiner's rejection of Claim 1 is respectfully requested.

With regard to Claim 2, it is respectfully submitted that the Malkin et al. patent does not disclose or suggest that the "harvesting software processes information contained in transmit hit/miss data and probability tables generated at the gateway", as is recited therein. The Examiner cited Figs. 2-3, column 7, lines 15-20 and lines 25-30, and columns 9-15

Column 7, lines 13-30 states

"The content server logic 268 includes a second statistics table 267, a second update statistics unit 269 for updating the second statistics table 267, and generating unit 266 for generating prefetch hint information. The second statistics table 267 stores information on requested information by the clients and potential prefetch information associated with the requested information. The statistics tables 261 contribute to the generation of prefetch hint information which his based on actual reference behavior of clients. Details on the statistics tables 261 are illustrated in FIG. 3B.
Operation of the second update statistics unit 269 is shown in greater detail with reference to FIG. 5. The generating unit 266 generates prefetch hint information from the information in the second statistics table 267 and the requested information. The generating unit 266 may be embodied as prefetch hint information (PHI) generation routine depicted in FIG. 6."

It is respectfully submitted that the Malkin et al. patent does not use the phrases "transmit hit/miss data" or "probability tables", and furthermore, there is clearly no disclosure or

suggestion that such transmit hit/miss data and probability tables are generated at the gateway, since there is no gateway disclosed or suggested in the Malkin et al. patent. The Examiner's conclusions are therefore based upon hindsight reconstruction.

Therefore, in view of the above and previously made arguments, it is respectfully submitted that Claim 2 is not obvious in view of the Malkin et al. patent and is patentable thereover. Dependent Claim 2 is also considered patentable based upon the patentability of Claim 1 from which it depends. Therefore, reversal of the Examiner's rejection of Claim 2 is respectfully requested.

Dependent Claims 3 and 4 are considered patentable based upon the patentability of Claim 1 from which they depend. Furthermore, the Examiner admitted that "Malkin does not explicitly disclose the limitation of NTTP objects", and has had to resort to the use of "Official Notice" that the concept and advantages of distributing NTTP objects to data networks is old and well known in the network communication art. It is respectfully submitted that this use of Official Notice is based upon hindsight reconstruction. Therefore, reversal of the Examiner's rejection of Claims 3 and 4 is respectfully requested.

With regard to Claim 5, it recites details of the gateway employed in the system of Claim 1. As was admitted by the Examiner, the Malkin et al. patent does not explicitly disclose a gateway. Therefore, it is respectfully submitted that the Malkin et al. patent does not disclose the details of the gateway recited in Claim 5, notwithstanding the use of "Official Notice" by the Examiner, which is based upon hindsight reconstruction.

It is respectfully submitted that the Malkin et al. patent does not disclose or suggest "a pseudo client for receiving an entitlement message indicating that that content has arrived at the gateway, for enabling the gateway as a sibling cache for the local cache, for requesting content to be transferred from the sibling cache to the local cache, for verifying that content has been transferred to the local cache during the transfer process, for disabling the gateway as a sibling cache of the local cache at the completion of the process" There is no discussion regarding clients, entitlement messages, or sibling caches disclosed or suggested in the Malkin et al. patent. The Examiner has admitted this.

However, the Examiner stated that "Malkin does disclose that the proxy server communicates with the local cache system of the client to request content on behalf of the client from the content server using the prefetch updating means and communication protocol used therebetween." It is respectfully submitted that this is not a disclosure or suggestion of the details of the gateway employed in the present system. The Examiner's arguments amount to unsupported speculation and hindsight reconstruction, and are unsupported by the teachings of the Malkin et al. patent.

Therefore, it is respectfully submitted that Claim 5 is not obvious in view of the Malkin et al. patent and is patentable thereover. Dependent Claim 5 is also considered patentable based upon the patentability of Claim 1 from which it depends. Therefore, reversal of the Examiner's rejection of Claim 5 is respectfully requested.

Dependent Claims 6 and 7 are considered patentable based upon the patentability of Claim 1 from which they depend. Therefore, reversal of the Examiner's rejection of Claims 6 and 7 is respectfully requested.

Claim 8 recites details of the harvesting software employed in the present invention. It is respectfully submitted that the Malkin et al. patent contains no disclosure regarding harvesting software that is coupled to a master cache and a gateway. It is also respectfully submitted that the Malkin et al. patent does not disclose or suggest "processes statistics derived from the master cache and the local caches to produces a list of content to add to the master cache and a list of content to delete from the master cache", "forms a pseudo client to retrieve and verify the content to be added to the master cache", and "transmits the verified content from the master cache to the local caches" as is recited therein. The term pseudo client is not even used in the Malkin et al. patent.

The Examiner stated that "Malkin discloses that proxy server logic 295 and content server logic 268 communicate statistical information to predictively retrieve or delete content based on probability calculated from clients request history." Even if it is assumed *arguendo* that the Examiner's statement is correct, this is not a disclosure or suggestion regarding forming a pseudo client. There is clearly no disclosure or suggestion in the Malkin et al. patent regarding forming a pseudo client, which the Examiner has admitted.

Therefore, it is respectfully submitted that Claim 8 is not obvious in view of the Malkin et al. patent and is patentable thereover. Dependent Claim 8 is also considered patentable based upon the patentability of Claim 1 from which it depends. Therefore, reversal of the Examiner's rejection of Claim 8 is respectfully requested.

Independent Claim 9 provides for a method for transferring content distributed by a data distribution system to a gateway into a local cache. The method comprises creating a pseudo client on the gateway, receiving an interrupt at the pseudo client indicating that that content has arrived at the gateway, enabling the gateway as a sibling cache for the local cache, requesting content to be transferred from the sibling cache to the local cache, verifying that content has been transferred to the local cache during the transfer process, disabling the gateway as a sibling cache of the local cache at the completion of the process, and causing the local cache to retrieve the content from the sibling cache until all content has been transferred.

It is respectfully submitted that there is no disclosure or suggestion in the Malkin et al. patent regarding "creating a pseudo client on the gateway". The term "pseudo client" is not used in the Malkin et al. patent nor is there any disclosure that would correspond to a pseudo client. There is no disclosure or suggestion in the Malkin et al. patent regarding "receiving an interrupt at the pseudo client". There is no disclosure or suggestion in the Malkin et al. patent regarding "enabling the gateway as a sibling cache for the local cache". There is no disclosure or suggestion in the Malkin et al. patent regarding "requesting content to be transferred from the sibling cache to the local cache". There is no disclosure or suggestion in the Malkin et al. patent regarding "verifying that content has been transferred to the local cache during the transfer

process" in the context of the recitation of Claim 9. There is no disclosure or suggestion in the Malkin et al. patent regarding "disabling the gateway as a sibling cache of the local cache at the completion of the process", There is no disclosure or suggestion in the Malkin et al. patent regarding "causing the local cache to retrieve the content from the sibling cache until all content has been transferred.

The detailed method steps recited in Claim 9 are considered patentable over the teachings of the Malkin et al. patent for the reasons argued above, and also with regard to Claims 1-8. Therefore, it is respectfully submitted that Claim 9 is not obvious in view of the Malkin et al. patent and is patentable thereover. Therefore, reversal of the Examiner's rejection of Claim 9 is respectfully requested.

Independent Claim 10 provides for a method for building a master cache used to transfer content by way of a data distribution system to a local cache. This method comprises processing statistics derived from the master cache and the local cache to produces a list of content to add to the master cache and a list of content to delete from the master cache, forming a pseudo client to retrieve and verify the content to be added to the master cache, and transmitting the verified content from the master cache to the local cache.

It is respectfully submitted that there is no disclosure or suggestion in the Malkin et al. patent regarding either "processing statistics derived from the master cache and the local cache to produces a list of content to add to the master cache and a list of content to delete from the master cache", or "forming a pseudo client to retrieve and verify the content to be added to the master cache". There is no disclosure or suggestion in the Malkin et al. patent regarding these steps. There is no pseudo client formed to retrieve and verify the content to be added to the master cache employed in the Malkin et al. system.

Therefore, it is respectfully submitted that Claim 10 is not obvious in view of the Malkin et al. patent and is patentable thereover. Therefore, reversal of the Examiner's rejection of Claim 10 is respectfully requested.

Dependent Claims 11 and 12 are considered patentable based upon the patentability of Claim 10 from which they depend. Also, with regard to Claim 11, the Malkin et al. patent discloses or suggests nothing regarding "processing information corresponding to probability distributions that the local caches satisfy requests from their respective users to predictively distribute the desired content to the respective users". Probability distributions are not discussed in the Malkin et al. patent. With regard to Claim 12, the Malkin et al. patent discloses or suggests nothing regarding "processing information contained in transmit hit/miss data and probability tables". Probability tables are not discussed in the Malkin et al. patent.

Therefore, it is respectfully submitted that Claims 11 and 12 are not obvious in view of the Malkin et al. patent and are patentable thereover. Therefore, reversal of the Examiner's rejection of Claims 11 and 12 is respectfully requested.

Claims 13-16 are considered patentable for the reasons argued above with regard to the other pending Claims. Therefore, it is respectfully submitted that Claims 13-16 are not obvious

in view of the Malkin et al. patent and are patentable thereover. Therefore, reversal of the Examiner's rejection of Claims 13-16 is respectfully requested.

Claims 17-19 address a specific satellite-based data distribution system which is implemented in a preferred embodiment of the present invention. Claims 17-19 are considered patentable for the reasons argued above with regard to the other pending Claims and because there is nothing disclosed or suggested in the Malkin et al. patent regarding anything other than a network (Internet) interconnecting the content and proxy servers, and nothing is disclosed or suggested regarding a satellite-based data distribution system. The Examiner's use of Official Notice to imply that a satellite-based data distribution system may be used in the Malkin et al. system clearly amounts to the use of hindsight reconstruction. This position is unsupported by the teachings of the Malkin et al. patent which is totally devoid of any concept of using a satellite-based data distribution system. Therefore, it is respectfully submitted that Claims 17-19 are not obvious in view of the Malkin et al. patent and are patentable thereover. Therefore, reversal of the Examiner's rejection of Claims 17-19 is respectfully requested.

In view of the above, it is respectfully submitted that Claims 1-19 are not obvious in view of the cited reference and are therefore patentable. Accordingly, it is respectfully submitted that the rejection of Claims 1-19 by the Examiner was erroneous, and reversal of the Examiner's decision is respectfully requested.

Respectfully submitted,

Kenneth W. Float
Registration No. 29,233

The Law Offices of Kenneth W. Float
Office address: 2 Shire, Coto de Caza, CA 92679
Mailing address: P. O. Box 80790, Rancho Santa Margarita, CA 92688
Telephone: (949) 459-5519
Facsimile: (949) 459-5520

APPENDIX

Claims 1-19 presented below are pending in this application.

1. A caching system for use with a data distribution system, comprising:

a master cache for receiving content for distribution by the data distribution system to one or more users;

a gateway for receiving content that is distributed by the data distribution system from
5   the master cache;

one or more local caches for storing the content received by the gateway destined for the one or more users; and

harvesting software coupled to the master cache and the gateway for processing information corresponding to probability distributions that the local caches satisfy requests
10   from their respective users to predictively distribute the desired content to the respective users.

2. The system recited in Claim 1 wherein the harvesting software processes information contained in transmit hit/miss data and probability tables generated at the gateway.

3. The system recited in Claim 1 wherein the content comprises http objects.

4. The system recited in Claim 1 wherein the content comprises nttp objects.

5. The system recited in Claim 1 wherein the gateway comprises:

a pseudo client for receiving an entitlement message indicating that that content has arrived at the gateway, for enabling the gateway as a sibling cache for the local cache, for requesting content to be transferred from the sibling cache to the local cache, for verifying that
5   content has been transferred to the local cache during the transfer process, for disabling the gateway as a sibling cache of the local cache at the completion of the process;

and wherein the local cache retrieves [56] the content from the sibling cache until all content has been transferred.

6. The system recited in Claim 5 wherein the Internet Protocol is used to communicate between the local cache and the sibling cache.

7. The system recited in Claim 5 wherein the Internet Cache Protocol is used to communicate between the local cache and the sibling cache.

8. The system recited in Claim 1 wherein the harvesting software:

processes statistics derived from the master cache and the local caches to produces a list of content to add to the master cache and a list of content to delete from the master cache;

forms a pseudo client to retrieve and verify the content to be added to the master cache;

5 and

transmits the verified content from the master cache to the local caches.

9. A method for transferring content distributed by a data distribution system to a gateway into a local cache, comprising the steps of:

creating a pseudo client on the gateway;

receiving an interrupt at the pseudo client indicating that that content has arrived at the

5 gateway;

enabling the gateway as a sibling cache for the local cache;

requesting content to be transferred from the sibling cache to the local cache;

verifying that content has been transferred to the local cache during the transfer process;

disabling the gateway as a sibling cache of the local cache at the completion of the

10 process; and

causing the local cache to retrieve the content from the sibling cache until all content has been transferred.

10. A method for building a master cache used to transfer content by way of a data distribution system to a local cache, comprising the steps of:

processing statistics derived from the master cache and the local cache to produces a list of content to add to the master cache and a list of content to delete from the master cache;

5 forming a pseudo client to retrieve and verify the content to be added to the master cache; and

transmitting the verified content from the master cache to the local cache.

11. The method recited in Claim 10 wherein the step of processing statistics comprises the step of processing information corresponding to probability distributions that the local caches satisfy requests from their respective users to predictively distribute the desired content to the respective users.

12. The method recited in Claim 10 wherein the step of processing statistics comprises the step of processing information contained in transmit hit/miss data and probability tables.

13. A caching system for use with a data distribution system, comprising:

a master cache for receiving content for distribution by the data distribution system to one or more user computers;

a gateway that is distinct from the one or more user computers for receiving content that
5   is distributed by the data distribution system from the master cache;

one or more local caches for storing the content received by the gateway destined for the
one or more user computers; and

harvesting software coupled to the master cache and the gateway for processing
information corresponding to probability distributions that the local caches satisfy requests
10   from their respective users to predictively distribute the desired content to the respective user
computers

14. The system recited in Claim 13 wherein the harvesting software processes
information contained in transmit hit/miss data and probability tables generated at the gateway.

15. The system recited in Claim 13 wherein the gateway comprises:

a pseudo client for receiving an entitlement message indicating that that content has
arrived at the gateway, for enabling the gateway as a sibling cache for the local cache, for
requesting content to be transferred from the sibling cache to the local cache, for verifying that
5   content has been transferred to the local cache during the transfer process, for disabling the
gateway as a sibling cache of the local cache at the completion of the process;

and wherein the local cache retrieves the content from the sibling cache until all content
has been transferred.

16. The system recited in Claim 1 wherein the harvesting software:

processes statistics derived from the master cache and the local caches to produces a list
of content to add to the master cache and a list of content to delete from the master cache;

forms a pseudo client to retrieve and verify the content to be added to the master cache;
5   and

transmits the verified content from the master cache to the local caches.

17. The system recited in Claim 1 wherein the data distribution system, comprises a
satellite-based data distribution system.

18. The method recited in Claim 9 wherein the data distribution system, comprises a
satellite-based data distribution system.

19. The system recited in Claim 13 wherein the data distribution system, comprises a
satellite-based data distribution system.